


# Anatomy of Kubernetes Operators for Databases

Peter Zaitsev / Sergey Pronin  
@ Percona



**HighLoad++**  
Becha 2021



# 1. Intro

The background features abstract, flowing shapes in shades of orange and red. On the left, there are several overlapping, semi-transparent orange shapes that curve upwards and to the right. On the right side, there are similar shapes in shades of red and pink, also curving upwards and to the left, creating a sense of movement and depth.

# Why Operators

- Customer-Driven Decision
- Containers=Kubernetes
- Operators = Simplification + Automation
  - Deployment
  - Management
- Percona Database-as-a-Service
  - On-prem or Cloud
  - Deploy on k8s

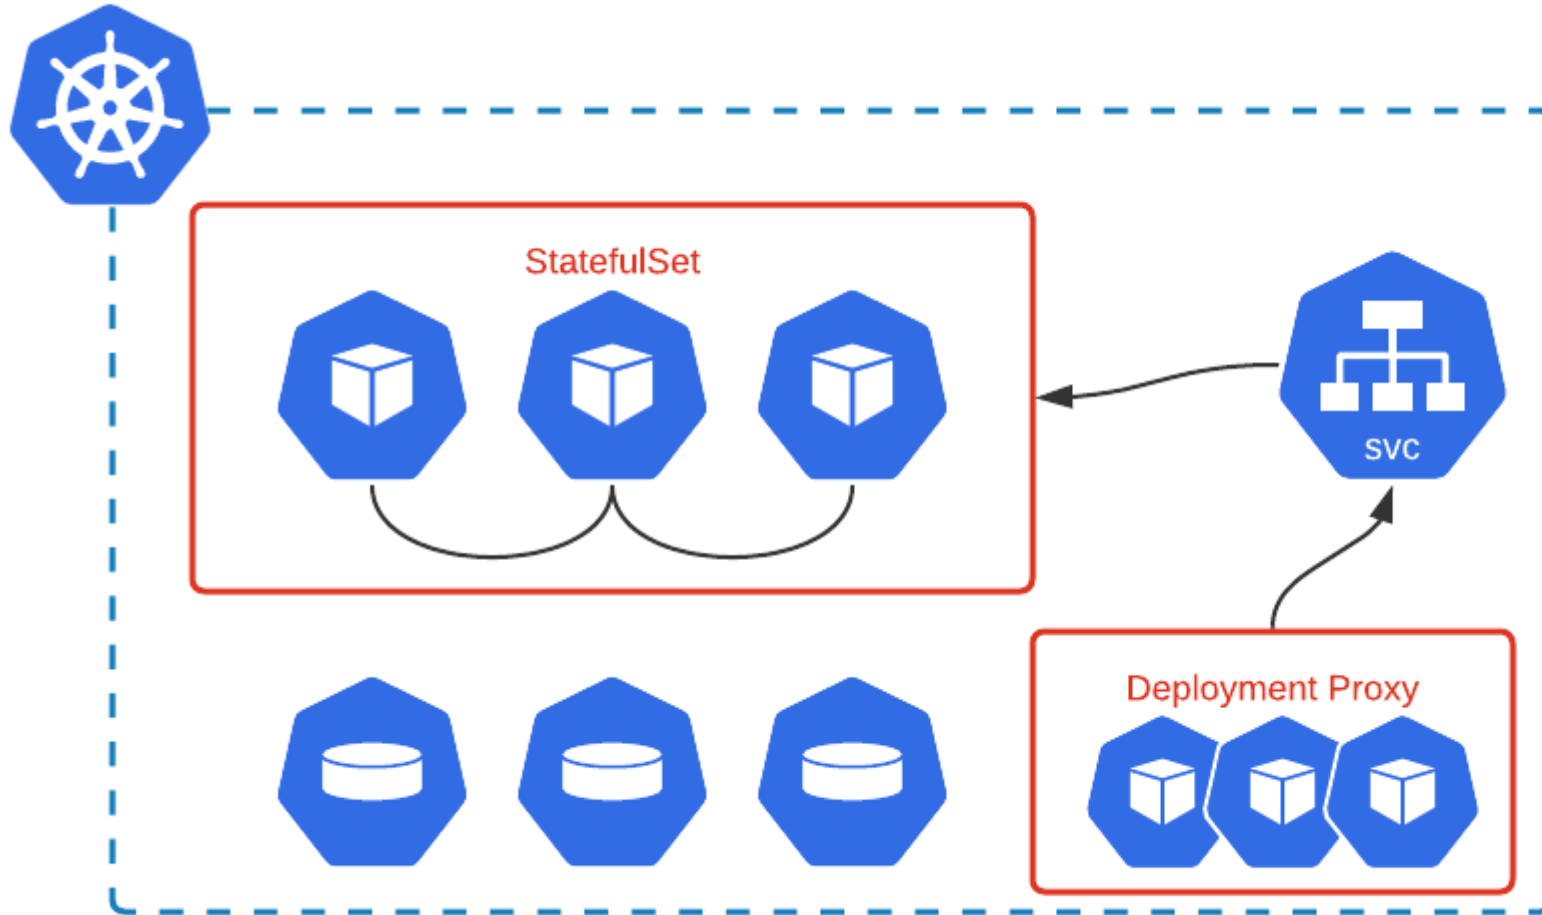
# Why Operators for DBs

- Cluster management is hard
  - K8S + database configuration is harder
  - Lots of things can go wrong
- Deployment is not enough, think operations
- It is not only DB
  - Proxy
  - Backups
  - Monitoring

## 2. The Hows


The background features abstract, flowing shapes in shades of orange and red. On the left, there are several overlapping, wavy shapes in various tones of orange. On the right, a large, sweeping shape in a vibrant red color rises from the bottom, partially overlapping the orange shapes. The overall effect is a modern, minimalist design with a warm color palette.

# Day 1 Operations - No Operator



1. Kubernetes
2. StatefulSet
  - a. DB Replication
  - b. Sustain
    - i. No downtime
    - ii. Automate recovery
3. Storage (PVC/Hostpath)
4. Service
  - a. Readers and writers
  - b. Query caching
5. Proxy
  - a. HA
  - b. Configure
6. Other
  - a. Customize DB
  - b. SSL, .....

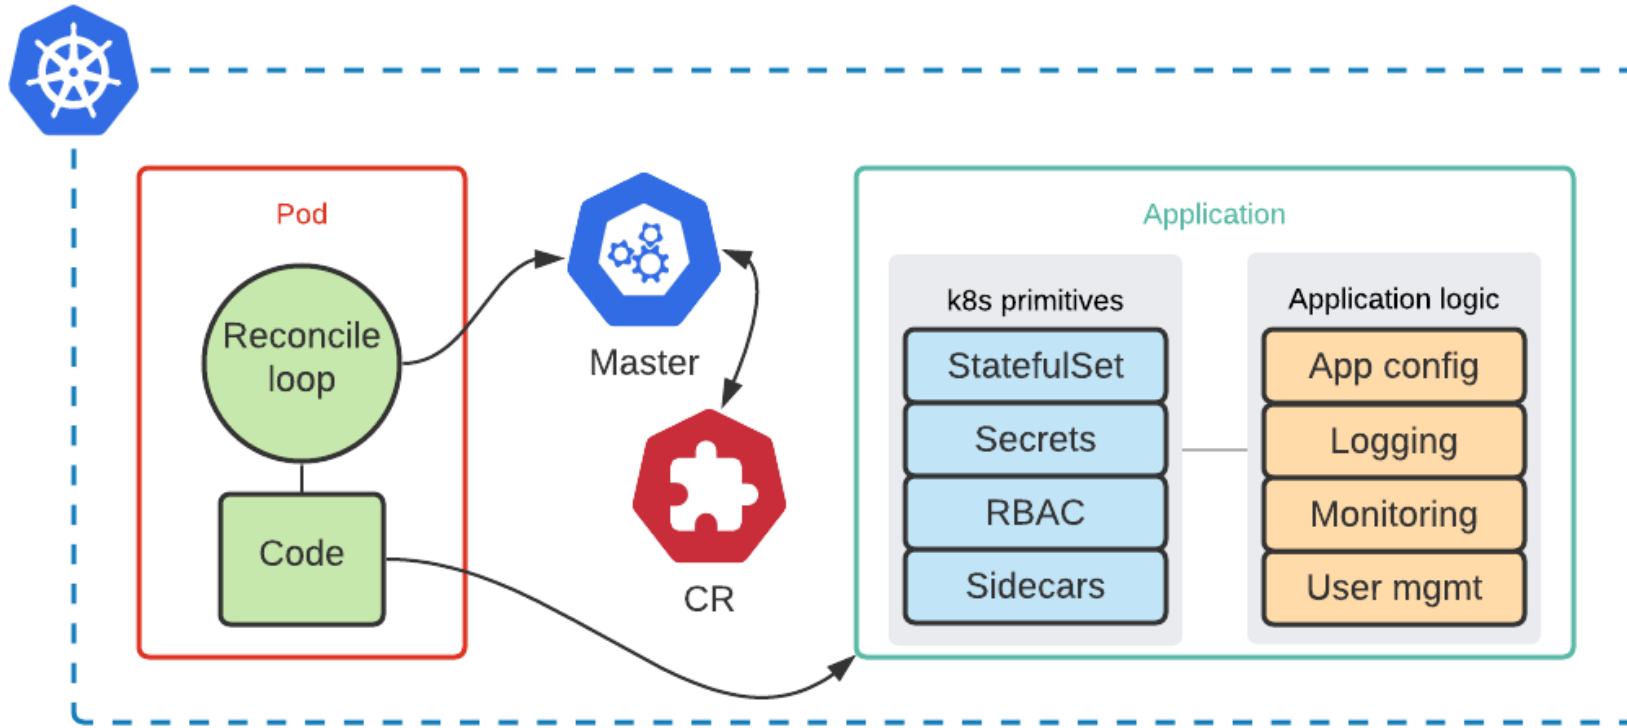
# Day 1 Operations - Operator



- Operator = Application + Kubernetes primitives
- Just give it the instructions

```
metadata:
  name: my-cluster
spec:
  pxc:
    size: 3
    image: percona/percona-xtradb-cluster:8.0.22-13.1
    affinity:
    annotation:
    resources:
    volumeSpec:
  haproxy:
    size: 3
    configuration:
    ...
```

# Operator - internals






# Day 2 Operations - Summary


- Reliability Engineering tasks
- Backup/Restore operations
- Scaling
- Upgrade
  - Operator
  - Database
  - Components (proxy, monitoring)
- DIY Database-as-a-Service

# Day 2 Operations - Scaling

## CPU/Mem




## Storage



# Day 2 Operations - backup / restore


## External backup



## Sidecar




## Custom Resource



# Day 2 Operations - Disaster Recovery

- Restore from backup
  - Without Custom Resource
- Replication
  - Kubernetes level
    - Availability Zones
    - Regional clusters
    - Federation
  - Data replication
    - Operator per k8s cluster



# Going multi-cloud

- Multi-cloud is a strategy
- Multi-cloud DBs is hard
  - Lock in on other services
  - Incompatibilities on DB engine/plugin level
- K8S does not change
  - Underlying infra does, but it is abstracted



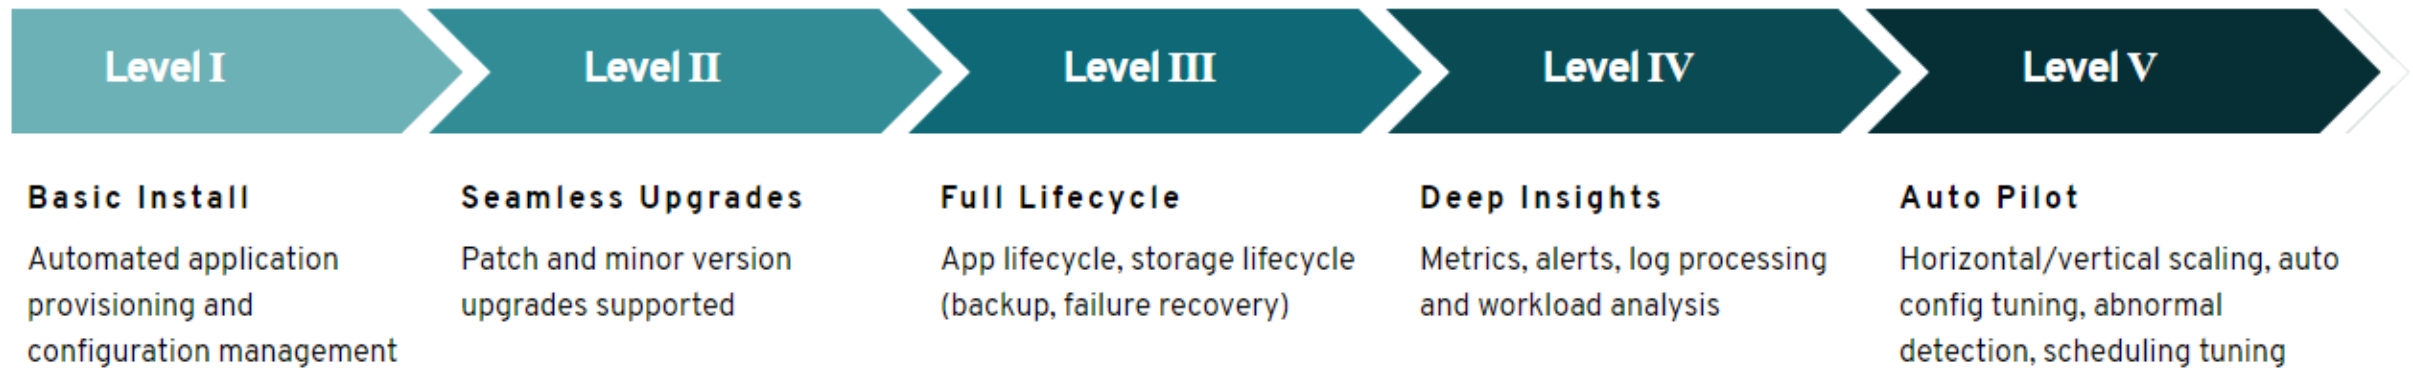
# Why not the cloud?

- Cost
- Vendor lock-in
- Multi-cloud
- Open source as a strategy

# 3. Summary



# Operators Capability Levels



Percona Operators check out on all levels, but not all items



# Summary

- Containers everywhere
- Databases on Kubernetes are hard
  - ~10 Pods, 20+ containers for one cluster
- Operators simplify management
  - Kubernetes primitives
  - Databases and components

# Plans for the Operators

- Feature parity
- Multi-cluster
- DBaaS integration
- PostgreSQL - polish and polish
- MySQL - v2 with async and Group Replication

Public Roadmap: <https://github.com/percona/roadmap/projects/1>

# Thanks!

Any Questions?

[pz@percona.com](mailto:pz@percona.com) / @PeterZaitsev

[sergey.pronin@percona.com](mailto:sergey.pronin@percona.com) / @sergeypronin